

15-418/618: Parallel Computer Architecture and Programming

Fall 2019 Syllabus

1 Course Details at a Glance

Lectures: MWF, 9:00am-10:20am, HOA 160
Instructors: Todd C. Mowry, GHC 9113, 268-3725, tcm@cs.cmu.edu
Brian P. Railing, GHC 6005, bpr@cs.cmu.edu
TAs: Abutalib Aghayev, agayev@cs.cmu.edu
Hanjun Li, hanjun1@andrew.cmu.edu
Joshua Kalapos, jkalapos@andrew.cmu.edu
Qidu He, qiduh@andrew.cmu.edu
Karen Bowman, kibowman@andrew.cmu.edu
Frank Fan, wfan@andrew.cmu.edu
Web Page: <http://www.cs.cmu.edu/418/>
Discussion: <http://www.piazza.com/cmu/fall2019/1541815618>

2 Course Description

From smart phones, to multi-core CPUs and GPUs, to the world's largest supercomputers and web sites, parallel processing is ubiquitous in modern computing. The goal of this course is to provide a deep understanding of the fundamental principles and engineering trade-offs involved in designing modern parallel computing systems as well as to teach parallel programming techniques necessary to effectively utilize these machines. Because writing good parallel programs requires an understanding of key machine performance characteristics, this course will cover both parallel hardware and software design.

Course themes include: (i) designing and writing parallel programs that scale effectively to large numbers of processors, (ii) understanding how parallel computers work (since this is very important in order to write fast and efficient parallel software), and (iii) thinking about efficiency (which is not equivalent to speed).

Here are the *learning objectives* for the course:

1. Design efficient and scalable parallel programs by applying algorithmic and architectural knowledge to utilize parallel computing resources;
2. Analyze and contrast the different parallel programming models (e.g., shared memory, heterogeneous, data-parallel, task-oriented, message-passing) and frameworks (e.g., OpenMP, Cilk, Cuda, MPI);
3. Investigate the common parallel and heterogeneous computing technologies and the costs associated with each;

4. Demonstrate appropriate measurement of parallel computing execution;
5. Understand the basics of ubiquitous computing workloads driving modern parallel systems design (e.g., graph processing, deep learning, imaging and graphics) and describe the models and frameworks best suited to each;
6. Explain in a clear language solutions to parallel computing problems.

3 Prerequisites

15-213 (*Intro to Computer Systems*) is a strict prerequisite for this course. We will build directly upon the material presented in 15-213, including memory hierarchies, memory management, basic networking, etc. While 18-447 (*Intro to Computer Architecture*) would be helpful for understanding the material in this course, it is not a prerequisite. Students are expected to be strong C/C++ programmers as there will be exposure to a variety of “C-like” parallel programming languages in this course. Previous experience with operating systems (15-410), distributed systems (15-440), or compilers (15-411) may prove to be very helpful.

4 Textbook

There is no required textbook for 15-418/618. If you are looking for additional (optional) reference material beyond the lecture notes, you might want to consider the following book:

- John L. Hennessy and David A. Patterson. *Computer Architecture, Sixth Edition: A Quantitative Approach*. Morgan Kaufmann, 2017

5 Course Work

Here is what you will be doing in this class.

Programming assignments: Students will complete four programming assignments. Assignment 1 will be performed individually. The remaining three assignments will be performed in groups of two.

Exams: There will be two in-class exams. The exams are closed-book, closed-notes. Each exam will cover roughly half of the course material (i.e. the second exam is not cumulative). There is no final exam in the course.

Final project: Over the last six weeks of the course, students will propose and complete a self-selected final project. By default, the final project will be performed in groups of two. Each team will present the project during a poster session (held during the normal final exam time) and produce a detailed write-up describing their work and results.

Class participation: To help students test their understanding of the class material and to prepare for the exams, we will have two different types of quizzes during the semester that are graded on a participation basis:

Mid-lecture quizzes: During most lectures, we will have a very brief quiz during the intermission period (roughly halfway through class) administered via Canvas. This quiz is usually a multiple choice question (or two), based upon the lecture material from that day. Students who complete at least 70% of these mid-lecture quizzes will receive full credit for this component of class participation.

Take home quizzes: To help students practice the types of questions that might appear on the exams, we will have take-home quizzes (between two and four over the semester). These quizzes must be completed individually without discussion with other students, but will be graded on a participation basis only.

5.1 Grading Policy

Your overall grade will be determined as follows:

Programming Assignments (4):	40% (exact weighting is TBD)
Exams (2):	30% (15% each)
Final Project:	25 %
Participation:	5 %

The grading cutoff points are: 90% (A), 80% (B), 70% (C), 60% (D). We will selectively consider raising individual grades for students just below the cutoffs based on factors such as attendance, class participation, improvement throughout the course, exam performance, project performance, and special circumstances.

Late hand-in policy. Each student is allotted a total of five late-day points for the semester. Late-day points are for use on programming assignments only. (They cannot be used for quizzes or final projects) Late-day points work as follows:

- A one-person team can extend a programming assignment deadline by one day using one point.
- A two-person team can extend a programming assignment deadline by one day using two points. (e.g., one point from each student)
- If a team does not have remaining late day points, late hand-ins will incur a 10% penalty per day (up to three days per assignment).
- No assignments will be accepted more than three days after the deadline.

5.2 Collaboration Policy

Students in 15-418/618 are absolutely encouraged to talk to each other, to the TAs, to the instructors, or to anyone else about course assignments. Any assistance, though, must be limited to discussion of the problems and sketching general approaches to a solution. Each programming project team must write their own code and produce their own writeup. **Consulting another student's or team's solution, or solutions from the internet, is prohibited. These and any other form of collaboration on assignments constitute cheating.** If you have any question about whether some activity would constitute cheating, just be cautious and ask the instructors before proceeding!

You may not supply code, assignment writeups, or exams you complete during 15-418/618 to other students in future instances of this course or make these items available (e.g., on the web) for

use in future instances of this course (just as you may not use work completed by students who've taken the course previously). Make sure to make repositories private if you use public source control hosts like github.

5.3 Student Wellness

Take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress.

All of us benefit from support during times of struggle. You are not alone. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is often helpful.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website at <http://www.cmu.edu/counseling/>. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night:

CaPS: 412-268-2922

Re:solve Crisis Network: 888-796-8226

If the situation is life threatening, call the police:

On campus: CMU Police: 412-268-2323

Off campus: 911

If you have questions about this or your coursework, please let us know.

6 Schedule

Table 1 shows the tentative schedule. The idea is to cover the lecture material in roughly the first $\frac{2}{3}$ of the semester (by meeting three rather than two days a week), so that you will have more time to devote to the class project in the last $\frac{1}{3}$ of the semester, and so that you can take advantage of all of the course lecture material in your projects.

Table 1: 15-418/15-618, Fall 2019, tentative schedule (subject to change).

Class	Date	Day	Topic	Assignments
1	8/26	Mon	Why Parallelism?	
2	8/28	Wed	A Modern Multi-Core Processor	
3	8/30	Fri	Parallel Programming Models	L1 Out
	9/2	Mon	<i>No Lecture: Labor Day</i>	
4	9/4	Wed	Parallel Programming Basics	
5	9/6	Fri	GPU Architecture and CUDA Programming	
6	9/9	Mon	Performance Optimization I	L1 Due, L2 Out
7	9/11	Wed	Performance Optimization II	
8	9/13	Fri	Parallel Application Case Studies	
9	9/16	Mon	Workload-Driven Performance Evaluation	
10	9/18	Wed	Snooping-Based Cache Coherence	
11	9/20	Fri	Directory-Based Cache Coherence	
12	9/23	Mon	Snooping-Based Multiprocessor Design	L2 Due, L3 Out
13	9/25	Wed	Performance Monitoring Tools	
14	9/27	Fri	Interconnection Networks	
15	9/30	Mon	Memory Consistency	
16	10/2	Wed	Scaling a Web Site	
	10/4	Fri	Exam I	
17	10/7	Mon	Implementing Synchronization	
18	10/9	Wed	Fine-Grained Sync, Lock-Free Programming	L3 Due, L4 Out
19	10/11	Fri	Transactional Memory	
20	10/14	Mon	Earthquake Simulation Case Study	
21	10/16	Wed	Heterogenous Paralelism, HW Specialization	
	10/18	Fri	<i>No Lecture: Fall Break</i>	
22	10/21	Mon	Tolerating Latency through Prefetching	
23	10/23	Wed	Domain-Specific Parallel Programming	L4 Due
	10/25	Fri	<i>No Lecture: Community Engagement Day</i>	
24	10/28	Mon	<i>Meetings to discuss project ideas</i>	
25	10/30	Wed	Addressing the Memory Wall	Project Proposal Due
26	11/1	Fri	Parallel Deep Neural Networks	
27	11/4	Mon	A Look Under the Hood	
	11/20	Wed	Exam II	
	11/22	Fri	<i>Meeting regarding Milestone Progress</i>	Project Milestone Due
	TBD		Project Poster Session (during Final Exam slot)	